

(56)

References Cited

PUBLICATIONS

Suganuma, "A Genetic Programming Approach to Designing Convolutional Neural Network Architectures", Cornell University, 2017 (Year: 2017).*

Bengio, "Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks", 2015 (Year: 2015).*

Abadi et al., "Tensorflow: A system for large-scale machine learning," In Proceeding of the 12th USENIX Conference, Nov. 2016, 21 pages.

Antoniou et al., "Data augmentation generative adversarial networks," arXiv, Mar. 21, 2018, 14 pages.

Bai et al., "Incorporating intra-class variance to fine-grained visual recognition," IEEE Xplore, arXiv, Mar. 1, 2017, 6 pages.

Baker et al., Designing neural network architectures using reinforcement learning, ICLR, arXiv, Mar. 22, 2017, 18 pages.

Bello et al., "Neural optimizer search with reinforcement learning," ICML, arXiv, Sep. 22, 2017, 12 pages.

Bergstra et al., "Random search for hyperparameter optimization," Journal of Machine Learning Research, Feb. 2012, 25 pages.

Brock et al., "Smash: one shot model architecture search through hypernetworks," arXiv, Aug. 17, 2017, 21 pages.

Ciregan et al., "Multi-column deep neural networks for image classification," IEEE Xplore, Jun. 2012, 8 pages.

Cubuk et al., "Intriguing properties of adversarial examples," arXiv, Nov. 8, 2017, 17 pages.

Deng et al., "Imagenet: A large-scale hierarchical image database," IEEE Xplore, Jun. 2009, 8 pages.

DeVries et al., "Dataset augmentation in feature space," arXiv, Feb. 17, 2017, 12 pages.

Devries et al., "Improved regularization of convolutional neural networks with cutouts," arXiv, Nov. 29, 2017, 8 pages.

Elsken et al., "Simple and efficient architecture search for convolutional neural network," ICLR, arXiv, Nov. 12, 2017, 14 pages.

Fei-Fei et al., "Learning generative visual models from few training examples: an incremental bayesian approach tested on 101 object categories," CiteSeerX, Aug. 2005, 12 pages.

Fukushima et al., "Neocognitron: a self organizing neural network model for a mechanism of visual pattern recognition," Springer-Verlag, 1980, 10 pages.

Gastaldi et al., "Shake-shake regularization," arXiv, May 23, 2017, 10 pages.

Goodfellow et al., "Generative adversarial networks," arXiv, Jun. 10, 2014, 9 pages.

Han et al., "Deep pyramidal residual networks," arXiv, Sep. 6, 2017, 9 pages.

He et al., "Deep residual learning for image recognition," arXiv, Dec. 10, 2015, 12 pages.

Hochreiter et al., "Long short term memory," Neural Computation, 1997, 32 pages.

Hu et al., "Squeeze-and-excitation networks," arXiv, May 16, 2019, 13 pages.

Inoue, "Data augmentation by pairing samples for image classification," arXiv, Apr. 11, 2018, 8 pages.

Jarrett et al., "What is the best multi-stage architecture for object recognition," IEEE Xplore, 2009, 8 pages.

Kornblith et al., "Do better imagenet models transfer better?" arXiv, Jun. 17, 2019, 24 pages.

Krause et al., "Collecting a large-scale dataset of fine-grained cars," Semantic Scholar, Jun. 2015, 2 pages.

Krizhevsky et al., "ImageNet Classification with Deep Convolutional Neural Network," Jun. 2017, 9 pages.

Krizhevsky, "Learning multiple layers of features from tiny images," Semantic Scholar, Apr. 2009, 60 pages.

Kumar et al., "Parallel architecture and Hyperparameter Search via Successive Halving and Classification," arXiv, May 25, 2018, 12 pages.

Laine et al., "Temporal ensembling for semi-supervised learning," arXiv, Mar. 15, 2017, 13 pages.

LeCun et al., "Gradient-based learning applied to document recognition," IEEE Xplore, Nov. 1998, 47 pages.

Lemley et al., "Smart augmentation learning an optimal data augmentation strategy," arXiv, Mar. 24, 2017, 12 pages.

Liu et al., "Hierarchical representation for efficient architecture search," arXiv, Feb. 22, 2018, 13 pages.

Liu et al., "Progressive neural architecture search," arXiv, Jul. 26, 2018, 20 pages.

Loshchilov et al., "SGDR: Stochastic gradient descent with warm restarts," arXiv, May 3, 2017, 16 pages.

Mahajan et al., "Exploring the limits of weakly supervised pretraining," arXiv, May 2, 2018, 23 pages.

Maji et al., "Fine-grained visual classification of aircraft," arXiv, Jun. 21, 2013, 6 pages.

Mania et al., "Simple random search provides a competitive approach to reinforcement learning," arXiv, Mar. 19, 2018, 22 pages.

Miyato et al., "Virtual adversarial training: a regularization method for supervised and semi-supervised learning," arXiv, Jun. 27, 2018, 16 pages.

Mun et al., "Generative adversarial network based acoustic scene training set augmentation and selection using svm hyper-plane," In Detection and Classification of Acoustic Scenes and Events Workshop, 2017, 5 pages.

Netzer et al., "Reading digits in natural images with unsupervised feature learning," NIPS, Jan. 2011, 9 pages.

Nilsback et al., "Automated flower classification over a large number of classes," IEEE Xplore, Dec. 2008, 8 pages.

Oliver et al., Realistic evaluation of deep semi-supervised learning algorithms, arXiv, Jun. 17, 2019, 19 pages.

Perez et al., "The effectiveness of data augmentation in image classification using deep learning," arXiv, Dec. 13, 2017, 8 pages.

Pham et al., "Efficient neural architecture search via parameter sharing," arXiv, Feb. 12, 2018, 11 pages.

Real et al., "Large-scale evolution of image classifiers," arXiv, Feb. 16, 2019, 16 pages.

Real et al., "Regularized evolution for image classifier architecture search," arXiv, Feb. 16, 2019, 16 pages.

Recht et al., "Do cifar-10 classifiers generalize to cifar-10?" arXiv, arXiv, Jun. 1, 2018, 25 pages.

Sajjadi et al., "Regularization with stochastic transformations and perturbations for deep semi-supervised learning," arXiv, Jun. 14, 2016, 9 pages.

Sato et al., "APAC: Augmented pattern classification with neural networks," arXiv, May 13, 2015, 9 pages.

Schulman et al., "Proximal policy optimization algorithms," arXiv, Aug. 28, 2017, 12 pages.

Simard et al., "Best practices for convolutional neural networks applied to visual document analysis," IEEE Xplore, Aug. 2003, 6 pages.

Simonyan et al., "Very deep convolutional network for large-scale image recognition," arXiv, Apr. 10, 2015, 14 pages.

Sixt et al., "RenderGAN: Generating realistic labeled data," arXiv, Jan. 12, 2017, 15 pages.

Sutskever et al [online], "Request for Research 2.0", Jan. 21, 2018, [retrieved Jul. 26, 2019], retrieved from URL<<https://openai.com/blog/requests-for-research-2/>>, 5 pages.

Szegedy et al., "Going deeper with convolutions" Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, 9 pages.

Szegedy et al., "Inception-v4, inception-resnet and the impact of residual connections on learning", AAAI, 2017, 7 pages.

Tarvainen et al., "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results" arXiv, Dec. 2017, 16 pages.

Tran et al., "A bayesian data augmentation approach for learning deep models" Advances in Neural Information Processing Systems, 2017, 10 pages.

Wan et al., "Regularization of neural networks using drop connection" International Conference on Machine Learning, 2013, 12 pages.

Xie et al., "Aggregated residual transformations for deep neural networks" arXiv, Apr. 2017, 10 pages.